

Atty. Dkt. No. 200311962-1REMARKS

Claims 1-31 were pending. Claims were 1-31 rejected. By the above amendments, Applicants have amended claims 6 and 8. Applicants hereby requests further consideration and re-examination in view of the amendments made above and remarks set forth below.

Drawing Objection and Specification Amendment:

The Office Action objects to the drawings because reference numeral 712 of figure 7A is not mentioned in the description. Reference numeral 712 refers to a decision block of a flow chart, where the decision block asks the question "Improvement technique?" and gives decision choices of "Yes" and "No." The description discusses figures 7A, 7B, and 7C, which together form a single flow chart, starting at page 25, line 15, through page 27, line 13. In particular, the originally filed paragraph starting at page 25, line 28, discusses "fourth or fifth steps, 708 or 710," which refer to decision blocks for "Ranking technique?" and "Threshold technique?" respectively, that immediately precede the decision block referred to by reference numeral 712 in the flow chart. The next paragraph starting at page 25, line 32, discusses "a seventh step 714." The originally filed paragraph starting at page 26, line 22, discusses the improvement technique, which is chosen by a yes answer to the decision block 712. From this context, it can be seen that reference numeral 712 refers to "a sixth step 712" for choosing or not choosing "an improvement technique." Accordingly, the specification paragraphs starting at page 25, line 28, and page 26, line 22, have been amended by the above amendments to include reference numeral 712 referring to the improvement technique. Thus, reference numeral 712 is now found in the description and there is no need to amend figure 7A. No new matter has been added by these amendments.

BEST AVAILABLE COPY

Atty. Dkt. No. 200311962-1

Specification Objections:

The Office Action objects to the originally filed title on the basis that the title of the invention is not descriptive. While the Applicants believe that the originally filed title is both descriptive and specific, the Applicants have amended the title of the invention to include the claim limitation of "instantiating a data placement heuristic" that is found in each of the independent claims.

The Office Action objects to related applications section regarding reference to co-filed related applications that are referred to by attorney docket number. By the above amendments, the paragraph starting at page 1, line 5, has been amended to replace reference to the co-filed related applications with their application numbers.

The Office Action requests that reference to the figures in the description of the drawings and the detailed description replace the word "figure" with the abbreviation "Fig." No authority is cited for this request and Applicants can find no authority that requires this. Accordingly, Applicants respectfully decline to make such amendments.

Additional Specification Amendments:

The detailed description has been amended to correct typographical errors of a "second step 204" and a "third step 206" with a "third step 206" and a "fourth step 208," respectively. Comparing flow diagram blocks 206 and 208 of figure 2 with the originally filed paragraphs shows that the language of the paragraphs refers to these blocks, respectively. Also, a "second step 204" is discussed earlier in the application and block 206 is preceded by two blocks in figure 2.

The detailed description has been amended to correct a typographical error which misidentified a method as "method 600," when from the context the reference is clearly to "method 700."

No new matter has been added by these amendments.

Claim Objections:

The Office Action objects to claims 6 and 8 for a lack of antecedent basis for the phrase, "the workload." Claims 6 and 8 have been amended to replace "the workload" with "a workload," which corrects the antecedent basis problem.

Atty. Dkt. No. 200311962-1Claim Rejections under 35 USC § 101:

The Office Action rejects claims 1-4 and 12-22 as being directed to non-statutory subject matter under 35 USC § 101. In particular, the Office Action rejects claims 1-4 and 12-22 as being directed to non-statutory subject matter because claims 1-4 and 12-22 do not produce a tangible result, which is respectfully traversed.

Claim 1 includes the claim limitation of *instantiating a data placement heuristic*. A data placement heuristic is a technique of placing data. Instantiating the data placement heuristic means to begin placing data according to the technique. The placement of the data may be a placement of data onto actual nodes or it may be a simulation of placement of the data onto nodes. For example, see the application at page 6, line 29, to page 7, line 3. In either case, the tangible result of placing data is achieved.

Claims 2-3 and 12-22 are dependent upon claim 1. Since claim 1 achieves a tangible result, claims 2-3 and 12-22 also achieve a tangible result.

Claim Rejections under 35 USC § 112:

The Office Action rejects claims 1-31 as being indefinite for the use of a relative phrase, which is respectfully traversed.

According to MPEP § 2173.02 at page 214, col. 1, last para., "The test for definiteness under 35 U.S.C. 112 ... is whether 'those skilled in the art would understand what is claimed when the claim is read in light of the specification,'" citing *Orthokinetics, Inc. v. Safety Travel Chairs, Inc.*, 806 F.2d 1565, 1576 (Fed.Cir 1986). According to MPEP § 2173.05(b) entitled, "Relative Terminology:"

The fact that claim language, including terms of degree, may not be precise, does not automatically render the claim indefinite under 35 U.S.C. 112, second paragraph. ... Acceptability of the claim language depends on whether one of ordinary skill in the art would understand what is claimed, in light of the specification.

The Office Action cites the phrase *within an allowable limit* that is found in each of the independent claims (claims 1, 23, 26, and 29-31) for finding that all of the claims (claims 1-31) are indefinite. The claim limitation in question reads *selecting a heuristic class which meets a performance requirement and which provides a replication cost that is within an allowable limit of a minimum replication cost*.

Atty. Dkt. No. 200311962-1

The particular phrase *within an allowable limit* within the claim language in question is definite for two reasons. First, the subject matter of the independent claims employs relative language because more specific language is not feasible. As taught in the description, the step of *selecting a heuristic class ... which provides a replication cost that is within an allowable limit of a minimum replication cost* may be a method of selecting a heuristic class. (See page 5, line 32, to page 15, line 22.) As taught in the description, determining general and specific lower bounds for the replication cost may employ one or more integer programs that are NP-hard, which means that there is no known algorithm that can solve the problem within any feasible time unless the problem size is small. (See page 5, line 32, to page 6, line 29.) Thus, determining an exact minimum replication cost either for the general problem or for a specific heuristic class is not feasible leading to the need for some relative language in the independent claims.

Second, the relative language is described in the description such that one of ordinary skill in the art would understand what is claimed when the claim is read in light of the specification. As taught in the description:

The allowable limit comprises a judgment made by an implementer depending upon such factors as a general lower bound (a lower general bound makes a larger allowable limit palatable), a cost of solving an additional specific integer program, and prior acceptable performance of the heuristic class modeled by the specific integer program. Typically, the implementer will be a system designer or system administrator who makes similar judgments as a matter of course in performing their tasks. (See page 15, lines 8-14.)

Thus, the relative language found in the independent claims, in light of the specification, is such that one of ordinary skill in the art would understand that what is claimed is a judgment made by an implementer such as a system designer or system administrator who makes similar judgments as a matter of course in performing their tasks. Therefore, claims 1-31 particularly point out and distinctly claim the invention.

Summary of Claim Rejections under 35 U SC § 102 and 35 USC § 103:

There are three sets of rejections under 35 USC §§ 102 and 103. In the first set of rejections, the Office Action rejects claims 1-9, 12-13, 16-18, 23-24, and 26-31 as being anticipated by Karlsson et al., *Do We Really Need Replica Placement Algorithms in Content Delivery Networks*, Tech. Report HPL-2002-220, Aug. 16, 2002 (referred to

Atty. Dkt. No. 200311962-1

herein as Karlsson I). In the second set of rejections, the Office Action rejects claims 1, 14-15, 19, and 20-22 as being anticipated by Karlsson et al., *A Framework for Evaluating Replica Placement Algorithms*, Tech. Report HPL-2002-219, Aug. 16, 2002 (referred to herein as Karlsson II). In the third set of rejections, the Office Action rejects claims 10-11 and 25 under 35 USC § 103 as obvious over Karlsson I in view of U.S. Pat. No. 6,466,980 to Lumelsky et al.

First Set of Claim Rejections under 35 USC § 102:

The Office Action rejects claims 1-9, 12-13, 16-18, 23-24, and 26-31 as anticipated under 35 USC § 102(b) by by Karlsson et al., *Do We Really Need Replica Placement Algorithms in Content Delivery Networks*, Tech. Report HPL-2002-220, Aug. 16, 2002 (referred to herein as Karlsson I), which is respectfully traversed.

Claim 1:

Claim 1 was rejected as anticipated by Karlsson I, which is respectfully traversed.

Claim 1 claims a method of determining data placement for a distributed storage system. A heuristic class is selected. The heuristic class meets a performance requirement and provides a replication cost that is within an allowable limit of a minimum replication cost. A data placement heuristic selected from a range of data placements according to the heuristic class is instantiated.

The Office Action rejection of claim 1 is incorrect in many respects. The Office Action refers to Karlsson I at page 3, col. 2, para. 4, to page 5, col. 1, para. 1, as teaching *selecting a heuristic class*, which is incorrect for at least two reasons. Karlsson I at the cited location is something on the order of 600 to 700 words, so it is being paraphrased here rather than being included here. To paraphrase, Karlsson I at the cited location discusses heuristics. But a heuristic is different from a heuristic class. According to the language of claim 1, a *heuristic class* has a *range of data placements* (i.e., each heuristic class covers a range of heuristics). Karlsson I teaches nothing about *a heuristic class* (i.e., a first reason that this citation to Karlsson I is wrong). And, Karlsson I teaches nothing about *selecting a heuristic class* (i.e., a second reason that this citation to Karlsson I is wrong).

Atty. Dkt. No. 200311962-1

The Office Action refers to Karlsson I at page 5, col. 1, para. 2, to col. 2, para. 1, as teaching *[a heuristic class] which meets a performance requirement*, which is incorrect for at least two reasons. As discussed directly above, the preceding citation to Karlsson I did not teach *a heuristic class* nor does this cited section teach *a heuristic class*. Therefore, Karlsson I does not teach *a heuristic class which meets a performance requirement*. And, consequently, Karlsson I does not teach *selecting a heuristic class which meets a performance requirement*.

The Office Action refers to Karlsson I at page 3, col. 2, para. 2, as teaching *[a heuristic class] ... which provides a replication cost that is within an allowable limit of a minimum replication cost*, with the reasoning that "cost is a constraint and the particular cost value at which the constraint is violated is the limit." The cited paragraph states:

An approximation method is the technique used to make the placement decision. The ones considered in this paper are *ranking*, *improvement*, and *Lagrangian relaxation*. Ranking starts with the computation of the cost impact of all possible combinations (within the metric scope) of placing one extra object on one node; sorts these costs and selects the best one that does not violate any constraints. If a constraint is violated, it tries the next placement in the list. This is repeated until no more objects can be placed. A *greedy ranking* heuristic recomputes the cost function after each object is placed. The specific improvement heuristic used in this paper is the 2-distance improvement heuristic [3]. It starts with an initial placement. This placement could be random or seeded by another heuristic. It then randomly picks one object and puts it on another node, making sure that the constraints are still satisfied. If this placement has a better cost it keeps it, otherwise it reverts back to the previous one. This process is then repeated a predefined number of times. Lagrangian relaxation [19] is a method that relaxes the constraints of the original problem by moving them into the cost function, which makes the new problem easier to solve.

A first reason that this citation is incorrect is that, as can be seen by reading the cited paragraph provided above, this paragraph is about computing costs according to the approximation methods of ranking, improvement, and Lagrangian relaxation not about *a replication cost that is within an allowable limit of a minimum replication cost*. And, second and third reasons that this citation is incorrect is that Karlsson I does not teach *a heuristic class* so it cannot teach *a heuristic class ... which provides a replication cost that is within an allowable limit of a minimum replication cost* and it cannot teach

Atty. Dkt. No. 200311962-1

selecting a heuristic class ... which provides a replication cost that is within an allowable limit of a minimum replication cost.

The Office Action refers to Karlsson I at page 5, col. 2, para. 2, to page 6, col: 1, para. 3, as teaching *instantiating a data placement heuristic selected from a range of data placement heuristics according to the heuristic class*, which is incorrect. The cited section of Karlsson states:

The placements of all the algorithms in this study were produced by *Coeus* [10]. This tool takes the problem definitions and heuristic primitives from Section 2.2 and produces the placement for that algorithm. In this way, it can run any previously published CDN algorithm and many variations thereof. We decided not to run any new algorithms, because we believe that the point of diminishing returns have already been reached for RPAs when used in the context of today's CDNs. Coeus uses existing algorithms and/or implementations of them when available. The placements produced by Coeus have been validated against existing implementations of algorithms when possible. All algorithms in this paper make decisions about placements using previously observed data, i.e. they do not know the future. The RPAs are run once every day unless otherwise specified, and we assume that they produce their placements instantaneously. The latter assumption has little impact on the results as the workload we study is quite stable.

Here are some details about the specific parameters of algorithms. The vicinity of the Hotspot algorithm is defined to be any client within 50 ms. Swap iterates for 5000 iterations and only generates feasible solutions. Lagrangian relaxation iterates for 500 iterations.

The decision costs of the algorithms are shown in Table 3. If the cost function is *dist* or *maxdist*, $K = 1$, as all objects are placed in the same way and only the placement of one object has to be calculated. We assume that the distance matrix and the fan-out vector change infrequently. Thus this information only needs to be fetched once. This can be seen in Table 3 for Ranking Dist and Fan-Out. The computation cost of Lagrangian relaxation is only valid for the so *readdist* problem formulation. To get the decision cost of a combined heuristic such as Greedy+, add the computation costs of the two heuristics and form the union of the number of messages cost and the union of the message sizes.

Karlsson I does not teach *the heuristic class* in this section or elsewhere. Therfore, Karlsson I does not teach *a data placement heuristic selected from a range of data placements according to the heuristic class* nor does Karlsson I teach *instantiating a data placement heuristic selected from a range of data placements according to the heuristic class*.

Atty. Dkt. No. 200311962-1

To anticipate a claim, a prior art reference must teach each and every limitation of the claim. Not only does Karlsson I not teach each and every limitation of claim 1; Karlsson I fails to teach many limitations of claim 1. Accordingly, claim 1 is allowable and an early allowance would be greatly appreciated.

Claims 2-8, 12-13, and 16-18:

Claims 2-8, 12-13, and 16-18 were rejected as anticipated by Karlsson I, which is respectfully traversed.

In the interest of brevity, the Office Action assertions regarding claims 2-8, 12-13, and 16-18 are not being individually addressed here. Rather, Applicants assert that claims 2-8, 12-13, and 16-18 are dependent upon independent claim 1. Dependent claims include all of the limitations of the claim upon which they depend. As explained above, claim 1 is not anticipated by Karlsson I. Thus, claims 2-8, 12-13, and 16-18 are also not anticipated by Karlsson I. Accordingly, claims 2-8, 12-13, and 16-18 are allowable and an early allowance would be greatly appreciated.

Claims 23, 24, and 26-31:

Claim 23, 24, and 26-31 were rejected as anticipated by Karlsson I, which is respectfully traversed.

In the interest of brevity, the Office Action assertions regarding claims 23, 24, and 26-31 are not being addressed here. Rather, Applicants assert that claims 23, 26, and 29-31 are independent claims that each include all of the limitations of claim 1 and that claims 24, 27, and 28 are dependent claims that each include all of the claim limitations of claim 1 by being dependent upon an independent claim that includes all of the claim limitations of claim 1. Since claim 1 is not anticipated by Karlsson I, claims 23, 24, and 26-31 are not anticipated by Karlsson I. Accordingly, claims 23, 24, and 26-31 are allowable and an early allowance would be greatly appreciated.

Second Set of Claim Rejections under 35 USC § 102:

The Office Action rejects claims 1, 14-15, 19, and 20-22 as being anticipated by Karlsson et al., *A Framework for Evaluating Replica Placement Algorithms*, Tech. Report HPL-2002-219, Aug. 16, 2002 (referred to herein as Karlsson II), which is respectfully traversed.

Atty. Dkt. No. 200311962-1

Claim 1:

Claim 1 was rejected as anticipated by Karlsson II, which is respectfully traversed.

Claim 1 claims a method of determining data placement for a distributed storage system. A heuristic class is selected. The heuristic class meets a performance requirement and provides a replication cost that is within an allowable limit of a minimum replication cost. A data placement heuristic selected from a range of data placements according to the heuristic class is instantiated.

The Office Action rejection of claim 1 is incorrect in many respects. The Office Action refers to Karlsson II at page 2, col. 1, para. 2, lines 6-11, as teaching *selecting a heuristic class*, which is incorrect for at least two reasons. Karlsson II at page 2, col. 1, para. 2, with underlining for lines 6-11 added here, states:

Section IV compares algorithms based on how "good" the produced placement is. A problem with the existing CDN literature is that algorithms are compared using minimized cost function values. Thus, comparison of algorithms with different problem definitions is impossible even when they target the same goal and system. One way to circumvent this problem is to compare their respective impact on the system performance or cost. In particular, we developed Coeus, an RPA generator that can produce the placement of most RPAs described in our framework to facilitate comparisons against the prior art. The power of this method is illustrated by comparing a number of algorithms never compared before, as they have diverse problem definitions. Finally, we present related work in Section V and conclude in Section VI.

This does not teach *a heuristic class* nor is *a heuristic class* taught anywhere in Karlsson II. While Karlsson II does employ the word heuristic (e.g., at page 1, col. 2, para. 2), a heuristic is different from a heuristic class. According to the language of claim 1, a *heuristic class* has a *range of data placements* (i.e., each heuristic class covers a range of heuristics). Further, Karlsson II teaches nothing about *selecting a heuristic class*.

The Office Action refers to Karlsson II at page 9, col. 2, para. 1, lines 2-3, as teaching *[a heuristic class] which meets a performance requirement*, which is incorrect. Karlsson II at page 9, col. 2, para. 1, with underlining of lines 2-3 added, states:

The performance metric to be used depends on what the system is used for. For a CDN, we use a client perceived latency threshold as the performance metric. The evaluation refers to the Cumulative Distribution

Atty. Dkt. No. 200311962-1

Function (CDF) of the latency to access objects, given the produced placements. The larger the ratio of access under the threshold, the better.

Merely mentioning a “performance metric” does not teach the limitation of *a heuristic class which meets a performance requirement*. As discussed directly above, Karlsson II does not teach *a heuristic class*. Therefore, Karlsson II does not teach *a heuristic class which meets a performance requirement*. And, consequently, Karlsson I does not teach *selecting a heuristic class which meets a performance requirement*.

The Office Action refers to Karlsson II at page 5, col. 1, para. 5, as teaching *[a heuristic class] ... which provides a replication cost that is within an allowable limit of a minimum replication cost*, which is incorrect. Karlsson II at page 5, col. 1, para. 5, states: “Fixed Threshold ($T(\text{threshold})$). An object is placed at a specific node, if the cost function is above or below a specified threshold. This approximation method is usually independent of the problem definition.” Merely mentioning a “cost function” does not teach the limitation of *a heuristic class ... which provides a replication cost that is within an allowable limit of a minimum replication cost*. As discussed above, Karlsson II does not discuss *a heuristic class*. Therefore, Karlsson II does not teach *a heuristic class ... which provides a replication cost that is within an allowable limit of a minimum replication cost*. And, consequently, Karlsson II does not teach *selecting a heuristic class ... which provides a replication cost that is within an allowable limit of a minimum replication cost*.

The Office Action refers to Karlsson II at page 2, col. 1, para., 2, lines 8-11, as teaching *instantiating a data placement heuristic selected from a range of data placement heuristics according to the heuristic class*, which is incorrect. Karlsson II at page 2, col. 1, para. 2, with underlining of lines 8-11 added, states:

Section IV compares algorithms based on how “good” the produced placement is. A problem with the existing CDN literature is that algorithms are compared using minimized cost function values. Thus, comparison of algorithms with different problem definitions is impossible even when they target the same goal and system. One way to circumvent this problem is to compare their respective impact on the system performance or cost. In particular, we developed Coeus, an RPA generator that can produce the placement of most RPAs described in our framework to facilitate comparisons against the prior art. The power of this method is illustrated by comparing a number of algorithms never compared before.

Atty. Dkt. No. 200311962-1

as they have diverse problem definitions. Finally, we present related work in Section V and conclude in Section VI.

Neither here nor elsewhere does Karlsson II teach *the heuristic class*. Therefore, Karlsson II does not teach *a data placement heuristic selected from a range of data placement heuristics according to the heuristic class*. And, consequently, Karlsson II does not teach teaching *instantiating a data placement heuristic selected from a range of data placement heuristics according to the heuristic class*.

To anticipate a claim, a prior art reference must teach each and every limitation of the claim. Not only does Karlsson II not teach each and every limitation of claim 1, Karlsson II fails to teach many limitations of claim 1. Accordingly, claim 1 is allowable and an early allowance would be greatly appreciated.

Claims 14-15, 19, and 20-22:

Claims 14-15, 19, and 20-22 were rejected as anticipated by Karlsson II, which is respectfully traversed.

In the interest of brevity, the Office Action assertions regarding claims 14-15, 19, and 20-22 are not being individually addressed here. Rather, Applicants assert that claims 14-15, 19, and 20-22 are dependent upon independent claim 1. Dependent claims include all of the limitations of the claim upon which they depend. As explained above, claim 1 is not anticipated by Karlsson II. Thus, claims 14-15, 19, and 20-22 are also not anticipated by Karlsson II. Accordingly, claims 14-15, 19, and 20-22 are allowable and an early allowance would be greatly appreciated.

Claim Rejections under 35 USC § 103:

The Office Action rejects claims 10-11 and 25 under 35 USC § 103 as obvious over Karlsson I in view of U.S. Pat. No. 6,466,980 to Lumelsky et al., which is respectfully traversed.

In the interest of brevity, the Office Action assertions regarding claims 10-11 are not being individually addressed here. Rather, Applicants assert that claims 10-11 are dependent upon independent claim 1. Dependent claims include all of the limitations of the claim upon which they depend. The Office Action asserts that Karlsson I teaches all of the limitations of claim 1 by asserting that Karlsson I teaches all of the limitations of claims 5, which is incorrect. As explained above, Karlsson I fails to teach many of the

Atty. Dkt. No. 200311962-1

limitations of claim 1. Therefore, Karlsson I does not teach many of the limitations of claims 10-11. Assuming for the sake of argument that there was a motivation or suggestion to combine Karlsson I and Lumelsky et al., such a combination would not teach or suggest these differences. Accordingly, claims 10-11 are allowable over Karlsson I in view of Lumelsky et al. and an allowance at an early date would be greatly appreciated.

Further, Applicants assert that claim 25 is dependent upon independent claim 23. Dependent claims include all of the limitations of the claim upon which they depend. The Office Action asserts that Karlsson I teaches all of the limitations of claim 23, which is incorrect. As explained above, Karlsson I fails to teach many of the limitations of claim 23. Therefore, Karlsson I does not teach many of the limitations of claim 25. Assuming for the sake of argument that there was a motivation or suggestion to combine Karlsson I and Lumelsky et al., such a combination would not teach or suggest these differences. Accordingly, claim 25 is allowable over Karlsson I in view of Lumelsky et al. and an allowance at an early date would be greatly appreciated.

Atty. Dkt. No. 200311962-1

Conclusion:

In view of the above, the Applicant submits that all of the pending claims are now allowable. Allowance at an early date would be greatly appreciated. Should any outstanding issues remain, the Examiner is encouraged to contact the undersigned at (408) 293-9000 so that any such issues can be expeditiously resolved.

Respectfully Submitted,

Dated: April 3, 2006



Derek J. Westberg (Reg. No. 40,872)

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- BLACK BORDERS**
- IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- FADED TEXT OR DRAWING**
- BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- SKEWED/SLANTED IMAGES**
- COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- GRAY SCALE DOCUMENTS**
- LINES OR MARKS ON ORIGINAL DOCUMENT**
- REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.